

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Using Analytics to Enhance Shelf-Space Management in a Food Retailer

Teresa Bianchi-Aguiar, Elsa Silva, Luis Guimarões, Maria Antónia Carravilla, José F. Oliveira

INESC TEC and Faculty of Engineering, University of Porto, Porto, Portugal,
{mtbaguiar@fe.up.pt, elsamsilva@fe.up.pt, guimaraes.luis@fe.up.pt, mac@fe.up.pt, jfo@fe.up.pt }

João Günther Amaral, Jorge Liz, Sérgio Lapela

Sonae MC, Maia, Portugal,
{jnamaral@sonaemc.com, jorgeliz@sonaemc.com, smlapela@sonaemc.com }

This paper describes the results of a collaboration with the leading Portuguese food retailer, addressing the shelf-space planning problem for allocating products to shop floor shelves. Our challenge was to bring analytical methods into the shelf-space planning process to improve the return on space and automate a process heavily dependent on the space managers' experience. This led to the creation of GAP, a decision support system that is today used daily by the space management team of the company. We developed a modular operations research approach that systematically applies tailor-made mathematical programming models combined with heuristics to derive the best layout of products on the shelves. GAP combines its analytical strength with the ability to incorporate different types of merchandising rules, to balance the trade-off between optimization and customization.

Key words: Retail operations, shelf-space allocation, mathematical programming, MIP-based heuristics

Sonae MC is one of the biggest Portuguese companies (ranked the 4th largest company in Portugal in 2014, with annual sales of 3.33 billion euros) operating a food retail business in Portugal. It is one of the core businesses of the Sonae Group, which also operates in other areas such as specialized retailing (sports goods, fashion, and electronics), shopping centers, and telecommunications.

Its *Continente* brand is the country's leading food retailer, and has been considered one of the most trusted brands in Portugal for the past 13 years. The company is a benchmark in the Portuguese market, having launched the country's first hypermarket in 1985. Today it has a network of 478 stores (with an additional 162 stores under franchising) covering the entire country. Stores adopt one of three major formats: *Continente Bom Dia*, convenience stores with an average sales

area of 800 m^2 (8,611 square feet); *Continente Modelo*, supermarkets located in medium-sized population centers, averaging 2,000 m^2 (21,528 square feet); and *Continente*, hypermarkets located in prime locations and offering an extensive and varied range of products and services with average sales areas of 9,000 m^2 (96,875 square feet). In total, Sonae MC has a sales area of 595,000 m^2 (6,404,527 square feet) and its strategy is to grow its convenience channel and to look for international growth opportunities.

Sonae MC is aware of the impact in-store planning has on customer satisfaction, sales effectiveness, and operations efficiency. In particular, it believes that a well-planned and innovative product organization on the shelves leads to higher visibility, consumer awareness, and increased demand for the products, as well as reduced inventory holding and handling costs. However, the short product life cycles, the increasing number of products available and the continually growing number of stores has led to an ongoing need to review shelf-space planning, making the process increasingly challenging for the company.

This paper describes the development, implementation and impact of an operations research (OR)-based approach to planning the allocation of products on the shelves. This is the result of a collaboration between the Information Systems and Innovation Department (ISI) and the Space Planning Department (SP) of Sonae MC, and a group of researchers from the Industrial Engineering and Management Department of the Faculty of Engineering of the University of Porto (FEUP).

Prior to this work, Sonae MC's shelf-space planning process was very time consuming, relying on an empirical use of space elasticities, lacking formal performance evaluation criteria and heavily dependent on the space managers' experience. The challenge consisted of incorporating analytical methods into the practice to automate the process, improve the return on space, and reduce stockouts and inventory costs, without disrupting (but somehow questioning) the company's policies. Guided by these objectives, three axes were defined for the project: Process Automation, Space Optimization, and Image Standardization.

We labeled the resulting decision support system as Automatic Generation of Planograms and use the acronym GAP. Space managers are using GAP daily to automatically generate shelf-space plans. GAP was developed on top of a modular architecture with innovative tailor-made mathematical programming models embedded in matheuristics that ensure its efficiency. In addition to the analytical advances, GAP offers the possibility of incorporating different types of allocation rules (also known as merchandising rules), allowing space managers to control the entire process, with a trade-off between customization with optimization.

The remainder of this paper is structured as follows. We start by describing how shelf-space is managed, first in Sonae MC and second from a more generic perspective, both from a practical and theoretical point of view. GAP is presented next, with a general discussion of its analytical approach and a description of the decision support system. We also discuss some details of the project development that were critical for its success. The impact of the project is carefully analyzed afterwards. We end with some brief concluding remarks, emphasizing that we are presenting a generic approach that is suitable for other retail companies. Note that this is a practice oriented paper and many details were omitted for the sake of simplicity. Additional papers will be referenced throughout the text for more technical details.

Shelf-Space Management at Sonae MC

The primary objective of retailing is to bridge the gap between the point of production and the point of sales, a goal that stresses the role of logistics and operations in this industry. Sonae MC has a centralized operations management department, responsible for planning all the operations for the stores nationwide. The space planning department, as its name implies, is responsible for managing the space available at the stores, an activity split into two main levels: a macro-space planning level that defines, on a long-term basis, the layout of the stores (divided by categories); and a shelf- (or micro-) space planning level that defines, for each category, the products' placement on the shelves. Shelf-space planning is a mid-term activity that updates shelf-space plans at an average rate of two to three times a year for more than 300 categories. This activity requires 23 full-time space managers.

The traditional shelf-space planning tool is a *planogram*, which is a virtual representation of the shelves, showing exactly where each product should be physically displayed and the inventory that it should hold. One planogram includes a variety of information that has to be carefully planned: the location of the products, the number of facings (visible items), the number of items stacked behind and above each facing, the packaging style, and the orientations (front, side, back, top), among others. Besides the most commonly used shelves, stores also have other fixture types such as chests, pallets and pegboards (bars with steel rods sticking out to hold peggable products like pens and pencils). Moreover, planograms are physically made of segments that are stacked together to form an aisle. Each segment has its own shelves: these can be placed vertically, aligned with the shelves of the other segments, or be placed with a different orientation, forming irregular planograms. Some examples are present in Figure 1.

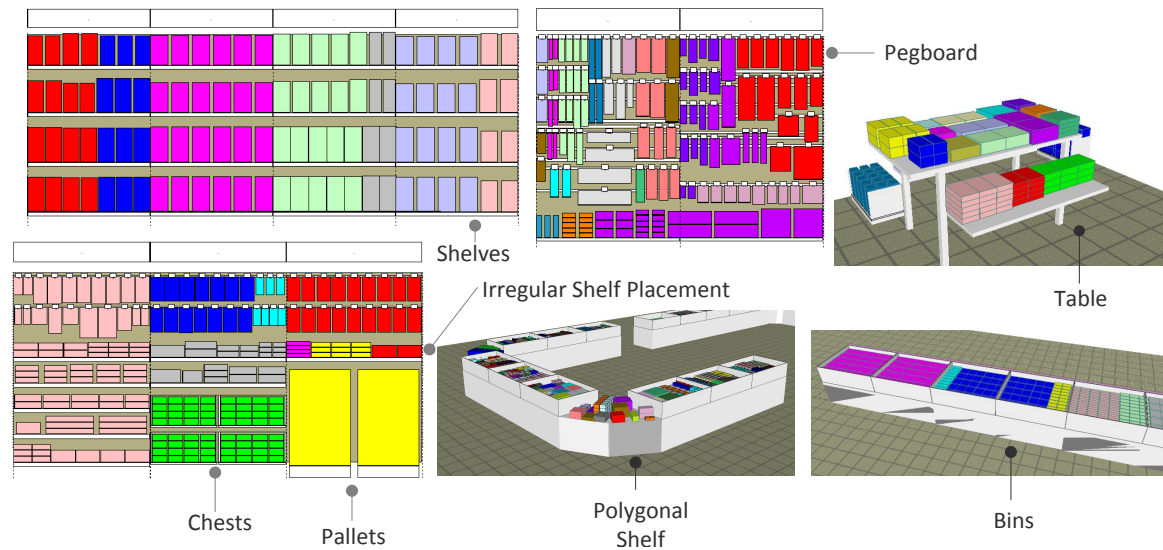


Figure 1 Planograms can have different types of fixtures. Some include more than one fixture type or present irregular shelf placements, resulting in irregular planograms.

At Sonae MC, planograms follow a complex structure of merchandising rules that try to reflect the consumer buying behavior and the strategy of the company (and of the suppliers) for the different categories. To this end, the company leverages its superior customer insight which it has gained from its successful loyalty card program, covering 3 out of 4 Portuguese households and linked to around 90 percent of sales. Moreover, the company maintains key partnerships with suppliers that have a deep knowledge of their categories, assuming the role of category captains. Space managers are also committed to developing planograms with a compelling visual look, devoting a great deal of effort to the task. Nevertheless, the attractiveness of the planogram is a subjective field and planograms depend on the space manager in charge. Figure 2 presents an example of a merchandising manual for a category, where we can see that products are usually grouped by families which are placed in rectangular shapes. Each planogram has a hierarchy of families that typically range from two to five criteria. For each criterion, the merchandising manual specifies the family type, the display orientation (either vertical or horizontal), the family precedences and, in some cases, additional information about preferred locations. Because of the strategic character of merchandising rules, this figure does not represent a real situation.

The process of updating shelf-space plans involves in-depth interaction with the commercial department which is responsible for managing categories. The process for a given category is as follows. The category manager (from the commercial department) triggers the process after specifying the product portfolio (assortment) for the stores, as well as the key merchandising rules

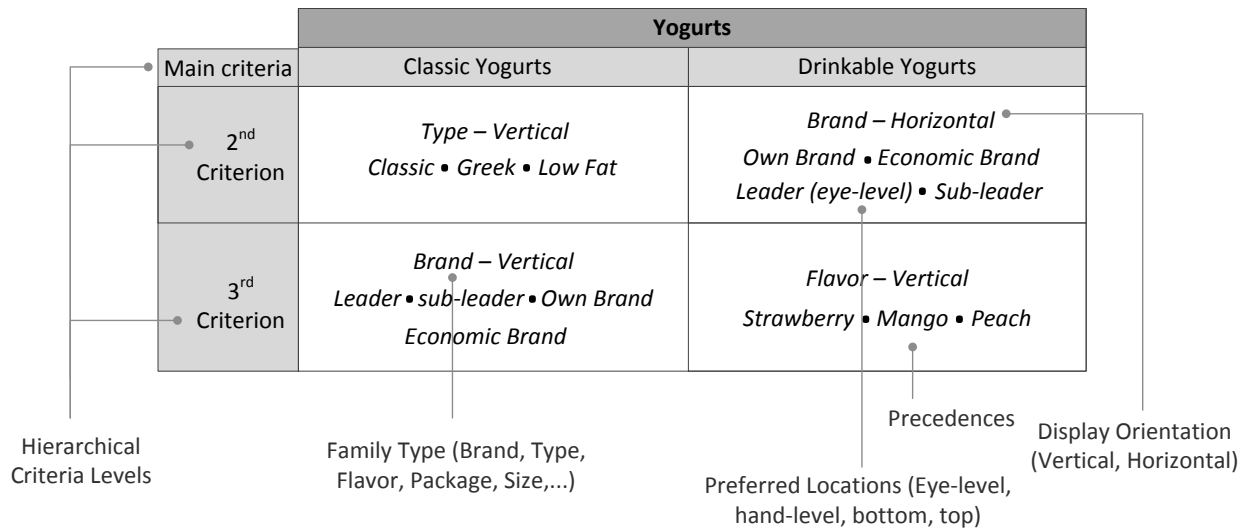


Figure 2 Merchandising rules reflect the strategy for implementing a given category (in this Figure the category is Yogurts). Each manual specifies between two and five hierarchical criteria levels with different types of information.

for their implementation. Product portfolios are not store-specific but are instead specified for clusters of stores with similar sales and space patterns to manage complexity and effort. Space managers start by generating a template planogram (known as a role planogram) for each cluster, where they carefully check how merchandising rules fit the space. In a collaborative effort between the space and category managers, the role planogram is then discussed and merchandising rules are tuned. Once validated by the category manager, it is then replicated for the remaining stores, by adjusting the product facings to the space of each store, while maintaining the same allocation rules. Figure 3 summarizes this shelf-space planning process, where the two key processes are highlighted: the Generation Process and the Replication Process. Note that the company has many categories to update and, at the beginning of each year, the category space planning processes are scheduled for the entire year.

During the shelf-space updating processes, space managers generate an average of 60,000 planograms each year. For this purpose, Sonae MC uses space planning software from one of the world's top three vendors, the JDA Software Group, Inc. This software has the necessary capabilities for creating and maintaining the planograms, including a space database with all the key information about the products and store equipment, and a visualization tool that provides realistic views of the shelves, the ability to easily handle products and powerful reporting tools. Although automatic tools for planogram generation are available in JDA's software, they do not accommodate all the inherent complexity of the merchandising rules. Therefore, space managers manually

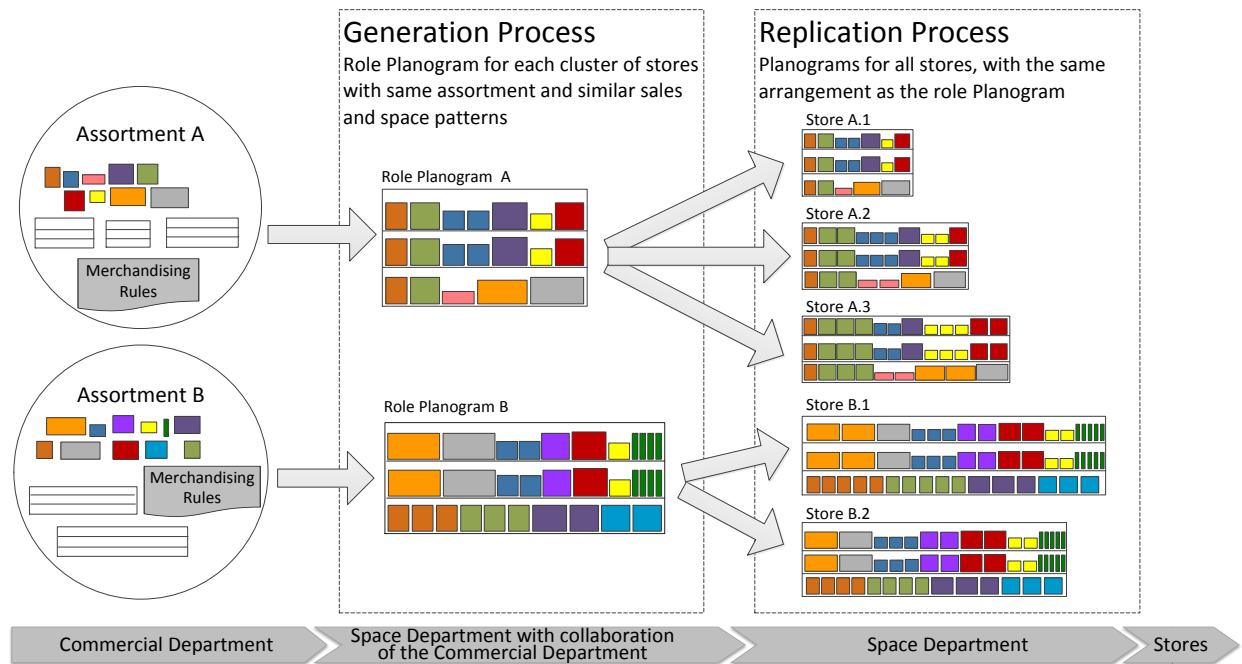


Figure 3 The micro-space planning process depends heavily on interaction with the commercial department and comprises two main processes: Generation and Replication.

developed their planograms by dragging and dropping the products onto the shelves, in a time-consuming activity that lasts on average three hours.

One of the most difficult challenges that we faced at the beginning of the project was the lack of formal criteria for evaluating planograms. Space managers were creating and evaluating planograms based on their intuition and personal judgment, as opposed to using analytical methods. Nevertheless, in most situations, they were empirically considering space elasticities and balancing the product days-supply values. When analyzing shelf inventory, days-supply is a common operational metric, measuring the number of demand days covered by the shelf stock. In balancing days-supply values, space managers were using a software highlighting tool that colored the products according to predefined days-supply intervals, and they sought to fit all the products within one interval. Moreover, some categories had alternative objectives such as meeting the brand market-shares.

In 2011, the stores went through a successful lean process that, among other things, changed their shelf replenishment strategy from a just-in-time policy (shelves were replenished frequently in small quantities during the day) to a single shelf replenishment operation each day, before the morning opening. This change of policy, and the fact that products normally have joint delivery

cycles from the central distribution centers (products delivered at the same days of the week), explains the reasoning behind balancing the days-supply values across the different products. By having all products supplied for a similar number of days, the number of shelf replenishment operations is reduced, the stock level for **long-tail (i.e. slow moving)** products is better controlled, stockouts for fast moving products are prevented, and it is also possible to reduce the backroom inventory (**store's warehouse**).

Sonae MC believed that analytics could help to improve shelf-space management, going beyond a simple planogram automation tool, providing the impetus for this project.

Theory and Practice of Shelf-Space Management

Most shoppers are susceptible to in-store marketing, mainly because of the low level of involvement that consumers have with in-store decisions. Additionally, reduced assortments and stockouts force them to search for substitute products, highlighting the role of space management. Experimental studies have been addressing the effect of space variables on the demand for products. These studies point to three main elasticities: *space elasticity* measures the increasing responsiveness of demand, as more space is allocated to a product, experiencing a decline in marginal returns at some point (Curhan 1972, Chandon et al. 2009); *location elasticity* highlights key display locations that help improve exposure, such as eye- or hand-level (Drèze et al. 1994); lastly, *cross elasticity* measures the interdependency between adjacent products and is assumed to be positive for complementary products and negative for substitute products (Corstjens and Doyle 1981). Additionally, the arrangement of products on the shelves can also have an important role in gaining consumers' attention. Thus, carefully organizing products in families can increase interest, while disorganization or excessive complexity (i.e. variations in the basic visual content) negatively affects the buying experience (Pieters et al. 2010).

According to a survey of U.S. retailers (Keltz and Sternecker 2009), the main drivers for space planning initiatives revolve around two main axes: maximizing selling space effectiveness, powered by the aforementioned effects, and tighter inventory control. However, the same survey concludes that the results are not meeting the expectations because “conventional assortment analytics and space tools do not deliver the optimization capabilities needed for success”. Software vendors mainly tackle the development of large-scale data processing technologies capable of addressing the complexity of shelf-space in practice, but with limited or no use of mathematical optimization, and a complete disregard for consumer demand effects. Therefore, automatically generated

planograms are still a mirage for most retailers and they often opt for generic planograms that fit clusters of stores.

Shelf-space management is an active field of research in retail operations management, where it is known as the Shelf-Space Allocation Problem (SSAP). Despite the practical relevance of the problem, scientific knowledge has been somehow misaligned with practice because most state-of-the-art mathematical models have strong limitations (Hübner and Kuhn 2012, Bai 2005).

The literature presents a great variety of models, mostly differing in their demand functions, which incorporate different estimates of some of consumer demand effects, ranging from complex multiplicative polynomial forms to simplistic linear profit functions. Nevertheless, most of these models have the common goal of maximizing demand by determining the product facings on each shelf, without considering their location within the shelves. The most relevant approaches to this work are from Corstjens and Doyle (1981), who were the first to present space elasticity in a polynomial form; Gajjar and Adil (2010), who propose a piecewise linearization to the space elastic demand function; and Yang and Chen (1999), who use an alternative model in the form of a linear multiple knapsack problem.

Perhaps the most important practical limitation from the aforementioned literature is that it neglects merchandising rules; more specifically, it disregards the existence of product families that specify associations between products on the shelves. Russell and Urban (2010) and Geismar et al. (2014) are the only authors who define the exact location of the products on the shelves and allocate the space to keep product families together, in uniform and rectangular shapes. Despite this, no model was able to reach an optimal solution for instances with more than 10 products.

Another key point is that the shelf-space allocation literature has given little focus to the cost side of the problem and most models do not explicitly consider inventory-related decisions. Two authors stand out in a more inventory-related stream: Baker and Urban (1988) presented the first model that considered the demand as a function of the instantaneous inventory level of an item, based on the economic order quantity (EOQ) model, and Urban (1998) proposed the first attempt to include shelf-space allocation in the inventory decision-making process. Nevertheless, these models are comprehensive and are only solved to optimality for a reduced number of products. The models also include practical limitations: they consider continuous shelf replenishment operations from the backroom and determine individual product replenishment policies. Our approach can also relate to this stream because we place special emphasis on inventory.

Finally, the literature regarding category captains is also interesting to this work, and a general treatment can be found in Kurtulus and Toktay (2009). Category captains are key suppliers that help retailers manage their categories, for example by consulting on the definition of merchandising rules. The use of external consultants with in-depth knowledge of the categories and the large number of factors behind merchandising rules reinforces the usefulness of rules as inputs to the space planning process.

Automatic Generation of Planograms (GAP): Overview

This section presents an overview of GAP. More precisely, we describe the new operations research-based process that we developed to help space managers at Sonae MC perform their daily activities more efficiently and effectively.

GAP's functionality centers around the generation of planograms that take into account the intrinsic complexity of the company's shelf-space planning process. In addition to its analytical advances, one of its most relevant features is its ability to allow users to control the entire planogram generation process, an activity that is highly dependent on the market, and so interconnected with the company's strategy. GAP can incorporate different types of merchandising rules that can change on the fly, allowing space managers to test different product layout strategies, and can be tuned to meet different performance evaluation criteria, such as the equilibrium of days-supply values across the products (the most typical planogram evaluation criteria in the company) or the maximization of expected demand based on space elasticity effects.

GAP has two major processes: *GAP Generation* and *GAP Replication*, responsible for each of the two processes identified in Figure 3: GAP Generation for generating planograms from scratch, following a set of merchandising rules, and GAP Replication for replicating a given planogram layout (possibly including non standard manual adjustments) to a new store. Although one may not initially see the relevance of GAP Replication, it was crucial to the success of the project because the large amount of fuzzy rules set by the commercial department makes the acceptance of fully automated planograms unlikely and inevitably leads to manual adjustments. Therefore, directly generating store specific planograms would break the validation point in the middle of the process and would result in the need for more validation and more manual adjustments, for both the category and space managers.

GAP consists of two main building blocks: (1) *GAP Optimizer* is the system's heart and contains all the analytical methods that make up GAP Generation and GAP Replication for generating

planograms; and (2) *GAP User Interface* provides all the tools for handling data and for managing the insertion of merchandising rules. The remainder of the section provides details on the analytical approach behind GAP Optimizer, the description of the overall decision support system, and an analysis of the key factors that led to a successful deployment in Sonae MC. Although GAP considers and integrates many types of fixtures (1), we will focus on the most common type: shelves.

Analytical Approach

Figure 4 presents the architecture of the GAP Optimizer. Both the GAP Generation and the GAP Replication processes were developed in a modular fashion, and systematically apply innovative tailor-made optimization models that were solved using mathematical programming-based heuristics (also known as matheuristics) to ensure fast solutions.

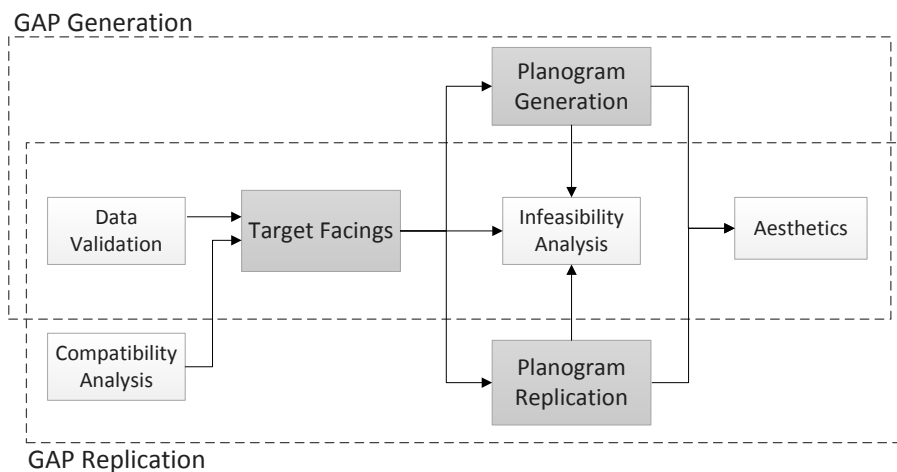


Figure 4 GAP Optimizer has a modular architecture and some modules are common to both GAP Generation and GAP Replication. The key modules are shaded.

Both the GAP Generation and GAP Replication processes start with a thorough data validation and compatibility analysis to guarantee that all the data are present and in accordance with the company's requirements. Once confirmed, the processes start by calculating the target number of facings for the products, with regard to one of the possible performance evaluation criteria. These values are then used as goals while generating the planograms, either in the planogram generation or planogram replication modules. Finally, because visual attractiveness is important, the last step is designed to improve the planogram aesthetics. If any of the shaded modules are unable to

generate valid solutions, the processes end with an infeasibility analysis to draw inferences on the causes.

We will now describe the key modules, with a high-level overview of the modeling and solution approaches included in the GAP Optimizer.

Target Facings The *Target Facings* module determines the number of facings that each product should have to maximize planogram performance, without considering any merchandising rules or other allocation constraints other than the shelf-space capacity. The reasons for estimating the target facings beforehand are twofold: it allows us to consider alternative performance evaluation criteria, and it reduces the complexity of the subsequent allocation problems. For brevity, this paper will focus on the company's most frequent planogram evaluation objective, which aims to balance the days-supply values across the products while considering space elastic demand.

At the center of the target facings calculation is the *space-to-sales* curve (Figure 5) which predicts the demand of a product as a function of the allocated shelf-space. This curve reflects the experimental findings regarding the space elasticity effect: the more space allocated to a product, the more consumer awareness the product has, leading to increased demand. Nevertheless, the marginal returns decrease as the shelf-space reaches a saturation point, resembling an “S” shape. We included a control parameter in the curve that specifies the maximum demand variation that can be explained by the shelf-space allocated to products, and consider demand forecasts (given as inputs) as the maximum value. Having captured the space elastic demand, we address the objective of balancing days-supply values by defining a set of days-supply intervals, and by limiting all the products to a unique interval. The possible days-supply intervals are calculated in a preprocessing phase using a user-defined interval length. These intervals are also depicted in the figure using dashed lines.

The target facings optimization model is formulated as a mixed-integer program (MIP) and embeds piecewise-linear approximations of each product's *space-to-sales* curve, obtained using the days-supply intervals. The model determines the target facings for each product that maximizes the planogram expected demand, subject to the shelf-space capacity, minimum and maximum number of facings and limited to the selection of a single days-supply interval. Additional details about this formulation are presented in Appendix A.

GAP Generation The *GAP Generation* module determines the products' placement on the shelves, subject to the user-defined merchandising rules. For each product, it determines the shelf (or shelves) where the product is to be allocated, its horizontal location within the shelves, and




Figure 5 The *space-to-sales* curve is at the center of the target facings calculation and predicts the demand of a product as a function of its shelf-space, limited to a maximum variation α . The vertical lines represent days-supply intervals.

the number of facings to be displayed. Therefore, the outcome is a fully defined planogram. This module aims to meet the target facings specified upstream (Target Facings module), while considering the location elasticity effects when choosing the products' placement. For that purpose, we have defined a set of *shelf attractiveness* curves that model the attractiveness of the shelves depending on their vertical locations (Figure 6). The shelf attractiveness may vary with the category and fixture type, which explains the alternative shapes.




Figure 6 The *Shelf attractiveness* curves model the attractiveness of the shelves depending on their vertical locations. We defined six curves (a-f).

Merchandising rules present a hierarchical structure of product families that are a key feature of every shelf-space plan. The products of each family must be placed together, in adjacent positions, and if a family spans more than one shelf, products have to maintain a continuous, uniform

and rectangular shape that can either be vertical (like columns, occupying the full height of the planogram) or horizontal (like lines, occupying the full length of the planogram). This hierarchical structure creates complex relationships between the products that highly constrain the solutions. We capture these relationships using a tree diagram, such as the one presented in Figure 7. The tree diagram starts with an initial node connecting the families from the first criteria, which define the first level of the tree. Each family is then connected to its downstream families, leading to a multilevel tree of product families.

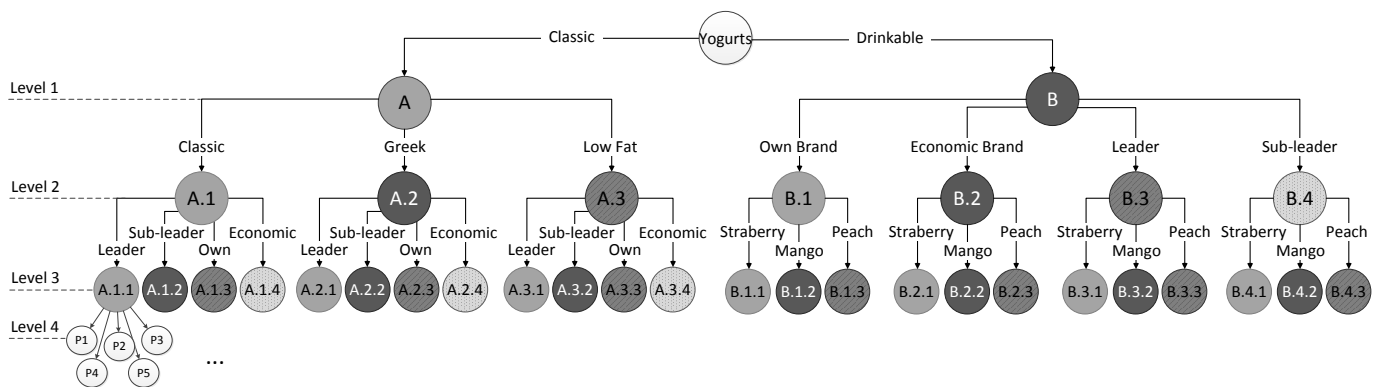


Figure 7 The tree diagram captures all the product family relations present in the merchandising rules. In this Figure the diagram reflects the merchandising rules from Yorguts, as present in Figure 2.

We formulated this problem using an innovative network flow MIP model. An intuitive network approach would associate one node to each product to find the sequence of products on each shelf. However, such an approach would lead to an extremely complex and intractable model because the traditional sequencing variables T_{ij} , stating whether each product i precedes or succeeds product j , would increase exponentially with the number of products. To address this, we explore the existence of a hierarchy in the product families and consider a set of multi-level family dependent networks. For each shelf, we start by defining a network with the families from the first level of the tree diagram. Afterwards, for each first-level family, we define a network with the corresponding downstream families and repeat the process until the last level is reached, with a set of disjoint product networks (see Figure 8 for a partial definition of the network resulting from the tree diagram in Figure 7). The model determines on the network sequences and then the overall product sequence is obtained by conveniently joining the product-level sequences. This approach also guarantees that the products belonging to the same family are consecutively placed on each

shelf, which is another requirement of the problem. Additional constraints ensure the coordination of each family between shelves for the rectangular shapes, as well as the display orientations. The precedence and special location rules, if any, correspond to variables that are fixed during a preprocessing step.

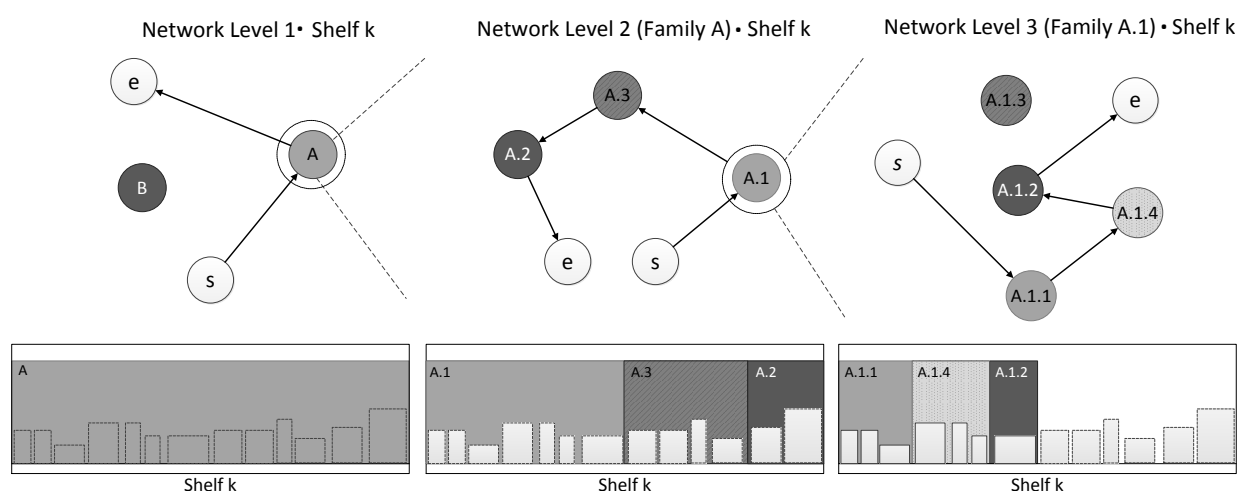


Figure 8 This Figure shows a partial representation of the multi-level family networks resulting from the tree diagram in Figure 7 (*s* and *e* represent the starting and ending node of each network).

One of the most relevant features of this formulation is its flexibility to accommodate any level of customization. Nevertheless, in less flexible scenarios, with highly detailed merchandising rules, the formulation will be more constrained, with many variables fixed, and optimization takes place with a reduced scope. For example, if the user chooses to set high-level family precedences, the remaining families are still sequenced by the model. In a more extreme case, if all sequences are defined, the formulation still determines the optimal shelf-space for each product.

When the problem size increases, it becomes intractable, even when resorting to the multi-level family networks. This limits the straightforward use of a commercial solver on the standard mathematical programming model. This challenge motivated the development of an approximate method. A mathematical programming based approach was chosen because it would be difficult to develop a highly customized heuristic considering all family-related merchandising rules, that would still be capable of generating high quality feasible solutions within reasonable time limits. Moreover, a MIP-based approach makes it possible to introduce new features in the problem with little or no effort. The formulation is then embedded in a metaheuristic that successively solves a

sequence of sub-problems exploring the hierarchy present in the product families. The matheuristic starts by allocating families from the first level and progressively moves down until reaching the product level. Technically, this approach is based on the relax-and-fix (R&F) framework: we consider the entire formulation in all iterations but families already considered in previous sub-problems have their variables fixed; families not yet considered have their variables relaxed to continuous values; and families from the current subproblem have integer variables. A backtracking scheme guarantees that whenever a sub-problem is infeasible, the heuristic shifts backward instead of forward, and solves a larger sub-problem by unfixing previous parts of the solution.

Table 1 presents a comparison of the performance of the matheuristic versus the straightforward use of a commercial solver on the model formulation for three real-world generation processes (artificially called CR_5 , AI_3 and BH_2), corresponding to a small, medium, and large instance, respectively. Both the solution quality and the running times proved its suitability for use in practice, especially because high quality solutions are obtained in less than one minute for the small and medium instances and less than five minutes for the larger one (in this case the commercial solver fails to delivery a feasible solution within one hour).

Table 1 We compared the performance of the matheuristic versus the straightforward use of a commercial solver on the standard mathematical programming model for planogram generation (time limit of one hour).

Name	Details				Formulation		Matheuristic	
	#Products	#Shelves	#Family levels	#Families	Optimality distance* (%)	Exec. Time**	Distance** (%)	Exec. Time***
CR_5	22	5	3	14	0.0 (OS)	00 : 00 : 07	0.0 (OS)	00 : 00 : 01
AI_3	41	9	5	27	8.5 (FS)	01 : 00 : 00	−0.1 (FS)	00 : 00 : 39
BH_2	131	5	4	26	− (IS)	01 : 00 : 00	− (FS)****	00 : 03 : 54

* Optimality distance (gap) - maximum deviation from the best integer solution to the upper bound;

** Distance - deviation from the solutions obtained with the matheuristic and a straightforward use of a commercial solver;

*** (hh:mm:ss); **** 1.8% Optimality gap, using the model's linear relaxation; (OS,FS,IS) - (Optimal, Feasible, Infeasible) Solution.

In summary, GAP Generation is particularly innovative because:

- It generates fully defined planograms which are uncommon in the shelf-space literature. Most models only tackle the space allocation to products disregarding product families, an essential feature in practice;
- The mathematical formulation includes modeling aspects that significantly improve its performance. The most important is the combination of commodity flow constraints for subtour elimination with the hierarchical networks of product families to reduce the combinatorial burden of product sequencing. Despite being introduced in this context, these modeling features can be explored in other sequencing problems having an hierarchical structure;

- It integrates many practical aspects of merchandising rules not considered before (multilevel product families, display orientation, product precedences) leading to more realistic problems in this area;
- The suggested matheuristic guarantees the generation of optimized planograms in practical-acceptable times even for large size instances with more than 200 products.

Both the formulation and the matheuristic are formally defined and compared in Bianchi-Aguiar et al. (2015b). An overview of the mathematical formulation is given in Appendix B.

GAP Replication Given a fully defined planogram (role planogram), the *GAP Replication* module reproduces a similar product placement for a new store without the need to provide merchandising rules or other types of reasoning behind the planogram construction. The new space is usually larger (in width, because planograms usually have the same height) but should have a similar shelf layout to ensure compatibility between the two planograms. Because the new store has a different demand pattern, the objective is to meet a new target number of facings, specified upstream and suitable for this store.

We formulated the replication problem as a MIP model. Although we mainly aim to adjust the product facings, the model necessarily determines the products' location, to guarantee that the new planogram fully complies with the role planogram. In particular, the following product placement information is considered:

- products are required to keep the same relative position as in the role planogram. In the case of shelves, this means that products maintain the same shelf level and they are placed following the same sequence;
- product families are required to keep their uniform rectangular shapes. The family continuity within each shelf is already ensured by keeping the same sequence. The rectangular shape is obtained by vertically aligning the first and the last products of the shape, which we call the left and right alignments. This allows us to consider shapes from the role planogram that may not necessarily be rectangular (resulting from manual adjustments).

One may say that the role planogram suffers a controlled “expansion” to keep all alignments. Note that during this process, we do not consider location effects on demand, as the relative product placement constrains such a decision.

Although solving the formulation in a commercial solver allows us to generate solutions within time limits that are acceptable in practice, we developed a second MIP-based heuristic to ensure

the scalability of the approach, especially for the planograms with irregular shelf placements (mis-alignments and interruptions in each shelf level), whose additional constraints greatly impacted the performance of the formulation.

Generically speaking, this matheuristic has three main steps. The first step generates an initial solution for the problem with the minimum display quantities for the products. The second step iteratively adds the remaining product facings to the planogram until no more space is available (or no more facings can be added). The last step tries to improve the solution by allowing the removal and insertion of new facings. Technically, this approach is an integration of two well-known MIP-based improvement heuristics: *fix-and-optimize* and *local branching*. Thus, in each iteration, we solve the model with some variables partially constrained in one of two ways: a subset of the variables is fixed to the values obtained in the incumbent solution (*fix-and-optimize*) or there is a limited number of changes allowed to the values obtained in the incumbent solution (*local branching*). One interesting aspect of this matheuristic is that it mimics the process followed by the space managers when manually replicating planograms.

Table 2 compares the results obtained using our solution approach and a commercial solver on the the model formulation. We use again three real-world instances ranging from small to large (E_1 , J_1 and G_1), to illustrate the efficiency and scalability of the approach.

Table 2 We compared the performance of the matheuristic versus the straightforward use of a commercial solver on the standard mathematical programming model for planogram replication (time limit of five minutes).

Name	Details			Formulation		Matheuristic	
	#Products	#Left alignments	#Right alignments	Optimality distance* (%)	Exec. Time	Distance** (%)	Exec. Time
E_1	49	3	4	0.0 (OS)	00 : 00 : 14	−0.1 (FS)	00 : 00 : 03
J_1	107	11	10	0.0 (OS)	00 : 01 : 27	−0.2 (FS)	00 : 00 : 02
G_1	240	32	35	0.12 (FS)	01 : 05 : 00	−0.2 (FS)	00 : 00 : 24

* Optimality distance (Gap) - maximum deviation from the best integer solution to the upper bound;

** Distance - deviation from the solutions obtained with the matheuristic and a straightforward use of a commercial solver (OS,FS,IS) - (Optimal, Feasible, Infeasible) Solution.

In summary, GAP Replication is particularly innovative because:

- We introduce the replication problem in the shelf-space literature, presenting a novel practice-oriented mathematical formulation. The replication process is done by most retailers which are forced to cluster their stores to manage categories efficiently. Therefore, we believe that this is an important step toward the use of analytic methods in shelf-space practice;

- For retailers which have complex merchandising rules, this is a very valuable tool allowing them to translate these hard to grasp and fuzzy rules in one example which is then the basis to create the remaining planograms;
- The three-phase mathematical programming based heuristic is able to solve most problems in less than three minutes guaranteeing the robustness and scalability of the approach. This is done by combining state-of-the-art programming heuristics such as *Fix-and-Optimize* and *Local Branching* in a single framework.

Both the formulation and the matheuristic are formally defined and compared in Bianchi-Aguiar et al. (2015a). To improve the comprehension of the current paper an overview of the mathematical formulation is given in Appendix C.

Aesthetics The formulations for planogram generation and replication focus on ensuring that the shapes are rectangular and disregard other aesthetic details, resulting in planograms with some display issues, such as large and irregular gaps between the products. The *Aesthetics* module is responsible for improving the attractiveness of the planogram and it considers two key factors for obtaining attractive displays: the way products are spaced throughout the planogram and whether the planogram is fully merchandised (i.e. full of facings). For that purpose, the generation or replication formulation (depending on whether it is a GAP Generation or Replication process) is executed again with all the decisions fixed to the incumbent solution, with the exception of the horizontal location of the products. The objective function is changed, first to minimize the empty space, and second to minimize the maximum spacing between two consecutive products when no more products fit the planogram. This latter objective distributes the empty space throughout the products.

Infeasibility Analysis Highly customized and detailed merchandising rules lead to significantly constrained generation and replication formulations which can compromise the existence of a feasible solution for the problems. Moreover, the target facings formulation can also be infeasible, which results in too many possible causes for the process ending without a valid solution. To overcome these data-related issues, we developed an *Infeasibility Analysis* module that searches for the possible infeasibility causes using a structured and logical procedure. It performs multiple runs of the infeasible formulation and, for each run, a problem feature or requirement is removed from the formulation. The process stops after identifying a source for the infeasibility (i.e. whenever the formulation is able to find a valid solution for the relaxed problem).

Decision Support System

GAP Optimizer requires the integration of different types of information obtained from multiple sources. If this information is not handled carefully, it may jeopardize the successful use of the application. To this end, another important building block is the GAP user interface that manages the data-handling process, executes the Replication and Generation processes with real-time status messages, and presents the generated planograms at the end, together with a full execution report. This report provides all the warnings and errors that occurred during the process and, when applicable, the infeasibility causes. That is, this interface is present throughout the entire process, helping in the liaison with the users. Figure 9 depicts the most relevant flows of information as well as snapshots from two interface forms: the project manager for handling the data and the generation manual for managing rules and configurations.

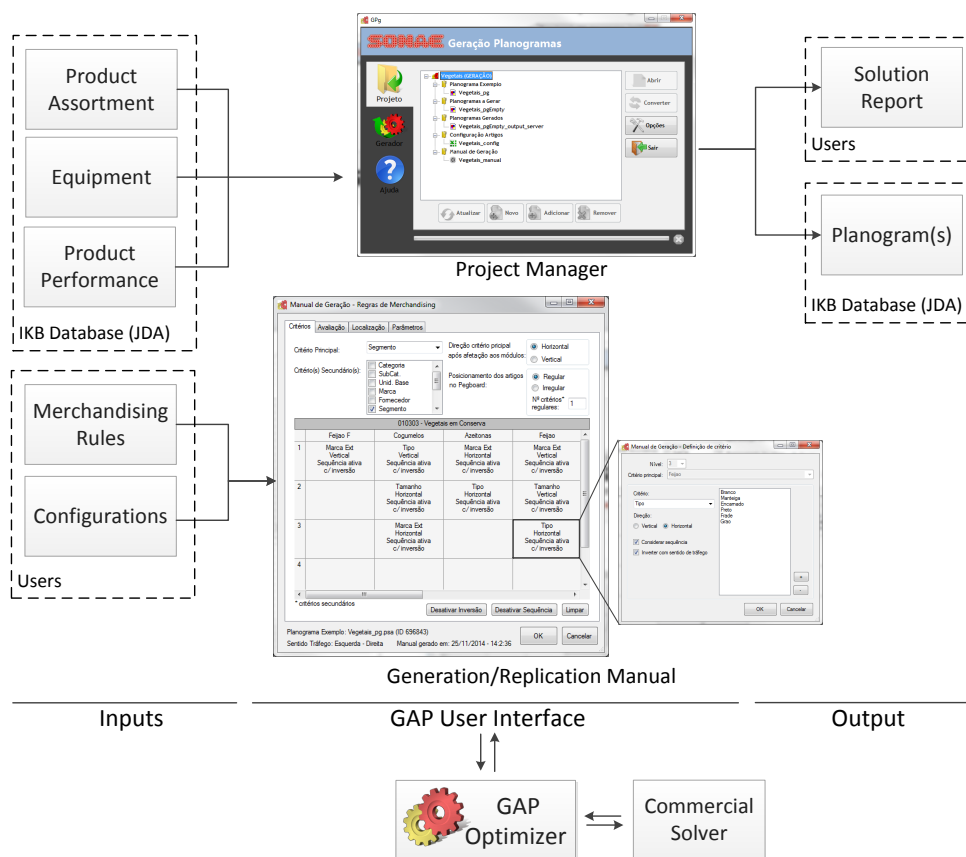


Figure 9 GAP has two main building blocks: GAP Optimizer and GAP User Interface. It requires the integration of information obtained from multiple sources.

The generation manual is one of the key parts of the overall system. Inspired by handmade merchandising manuals (c.f. Figure 2), this form presents a familiar interface (to space managers)

for the configuration of merchandising rules adopting a intuitive approach. Space managers have a high degree of flexibility in defining these rules, and in each run they can choose the level of customization that they require in the generated planograms. For the advanced users, several other configurations are available, from alternative location-elasticity curves and planogram evaluation criteria to control parameters and tolerances. Each generation manual can be saved, consulted and reused in multiple processes. Most importantly, it can evolve as the space managers evaluate planogram solutions and realize possible changes to the planograms.

The GAP Optimizer is a C++ program with all the models embedded in the code. The company acquired a commercial solver and the formulations are executed using a C++ library from the solver. The GAP user interface is developed using Windows Forms, and all the communications between the two building blocks use XML files. Both the GAP Optimizer and the GAP user interface are executed on a dedicated server and all space managers have access to the interface using a remote desktop connection in a terminal-server architecture. At the moment, GAP does not have a direct connection to the space database and the information is manually exported and imported to the interface. Given the success of the project, Sonae MC is now studying more efficient infrastructures, both for communicating with the space planning database and for accessing the server from the space managers' terminals.

Project Development

The project was kicked-off in March 2012 and ran until July 2014. The two processes, GAP Generation and GAP Replication, were developed sequentially and each involved three main phases: *Requirements Definition*, *Prototype & Proof-of-concept* and *Testing & Validation*. From the organizational standpoint, it included a team from FEUP, responsible for the development of GAP, and two Sonae MC teams: a team from the Space Department, responsible for validating requirements and testing GAP, and a team from ISI (the Information Systems and Innovation Department), responsible for integrating GAP into the information systems of the company.

We strongly believe that there were some key factors associated with GAP's design and project management which were crucial to the project's success. The decision to divide GAP in two processes played a vital role, from the perspective of both the space and commercial department, because GAP's development did not disrupt current practices. Starting the implementation with *GAP Replication* was also in hindsight a wise decision for two reasons. First, the replication process was quicker to implement and provided more consensual solutions, helping to bring the space

managers on board with GAP at an earlier stage. Second, it allowed us to deepen our knowledge of the complex structure defining the merchandising rules, which was vital for *GAP Generation*.

Another key aspect was the close collaboration that occurred between the three space managers who were part of the project team, whose role was essential from requirements gathering and problem definition to testing and validation phases. Weekly meetings between FEUP and the space managers were important milestones for validating new developments. This continuous process allowed us great flexibility in implementing GAP and led us to develop an application tailored to Sonae MC's requirements. Additionally, three space managers tested and validated GAP using different categories of products, which was also significant for building (and communicating) their internal confidence in the application.

Nevertheless, we faced some challenges in this collaboration between academic researchers and industry practitioners who have different objectives, incentives and time horizons. In particular, at the beginning of the project we found that the space managers were highly resistant to adopting the new process. This was gradually overcome as we attempted to keep them updated on the evolution of the project, which improved their commitment to GAP and allowed them to understand the potential of the application.

Training sessions were also given before the roll-out of each of the two processes and, given the systems' complexity, they were crucial to engaging space managers. These sessions included the analysis of planograms with unexpected characteristics obtained while using GAP and the development of a check-list for systematically looking for alternative solutions in these situations. Moreover, the execution report also played a vital role in allaying the concerns of space managers when GAP produced an unexpected solution and, more importantly, when it did generate one.

Last – and perhaps one of the most important key factors for the success of the GAP implementation – the space planning and innovation directors showed enormous commitment and provided their full sponsorship during the whole project.

Impact

Today, GAP is used daily by the entire micro-space team at Sonae MC, both for the generation and replication of planograms. This section describes how it enhanced shelf-space management for each of the three axes that were identified for the project: Process Automation, Space Optimization and Image Standardization.

Automation: From Planogram Construction to Planogram Evaluation

Perhaps the first and most straightforward impact of the project was on the space management process, which led to better processing times and a positive change of paradigm.

During the first months after the roll-out, all space managers were encouraged to use GAP in their daily tasks, and to register the number of planograms that GAP could generate, as well as the process execution times (including the duration of data handling, creation of the generation and replication manuals, and handmade adjustments to the final solution). The execution times were later compared with the legacy process and the results were encouraging: based on an analysis of over 400 planograms in eight different categories, space managers were able to automatically generate 8 percent of the planograms and the category space management processes took on average 48 percent less time (46 percent less in the generation processes and 50 percent less in the replication processes). Moreover, space managers also highlighted that these time reductions could be more significant in the future, after they become more familiar with the new software and totally, or partially, reuse the generation manuals that they carefully developed for the first processes.

Additionally, GAP shifted the space managers' focus from planogram construction to planogram evaluation, allowing them to concentrate on additional activities, such as market trend studies and experiments with alternative merchandising rules. Therefore, this change of paradigm brought increased responsibility to space managers, and gave them an analytical tool to support their decisions during the meetings with the commercial department.

Optimization: Targeting Optimality for All Customization Levels

The impact of GAP Generation and GAP Replication on optimization needs to be measured independently.

One of the greatest advantages of GAP Generation is its flexibility to generate either highly customized solutions or more demand driven (and concurrently innovative) solutions, based on the specifications in the generation manual. Figure 10 depicts this flexibility by showing three planograms: the first was handmade by a space manager and the remaining two were generated with GAP, first using a fully defined generation manual (high customization) and then using the same manual but removing family precedences and display directions (low customization). While the first generated planogram is largely similar to the handmade planogram, the second presents an alternative reasoning behind its creation where the intention was to put the most popular families in the premium vertical and horizontal locations.

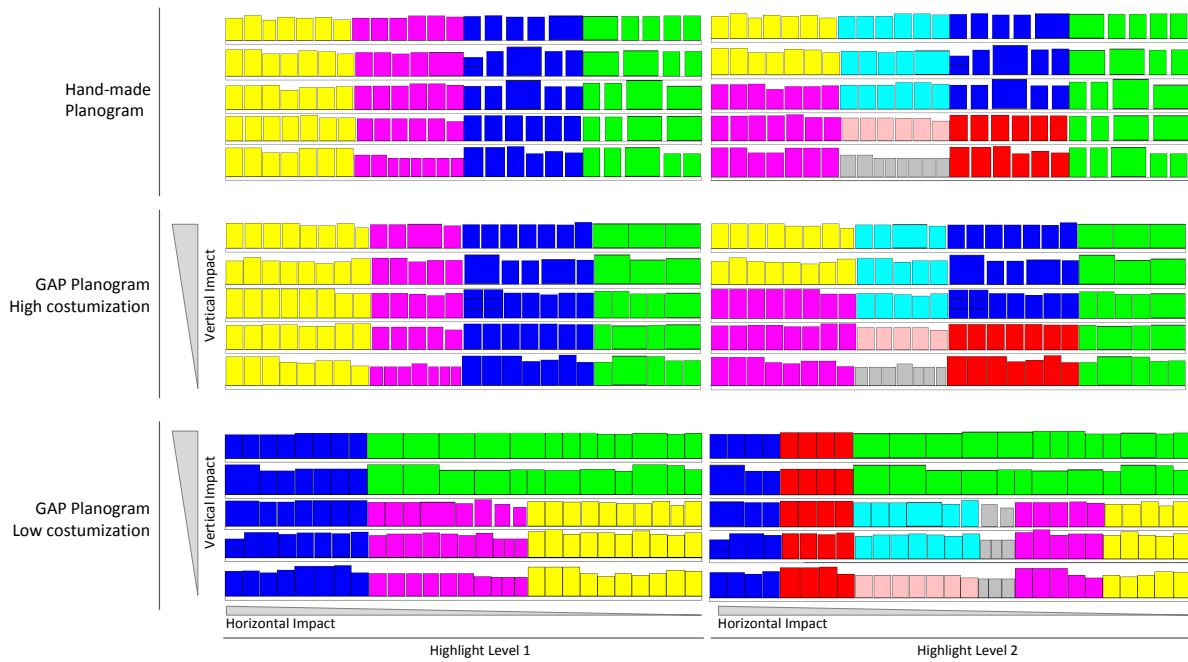


Figure 10 This Figure shows two planograms generated with GAP using two different levels of customization, and their comparison with a handmade planogram.

Regardless of the customization level, GAP Generation always uses the available degrees of freedom to optimize the number of product facings and the products' location. To assess the impact of GAP Generation on the planograms' performance, we analyzed 25 generation processes (one of which was the example described above). These examples were carefully selected during the proof-of-concept phase to guarantee that all specificities of the categories were covered. The impact was evaluated by measuring four performance metrics: potential sales increase (estimated using the location elasticity curves with a maximum impact of 20%); days-supply balance (measured in terms of the average and the standard deviation reductions); and planogram filling rate (defined by the ratio of the linear space utilized and the overall available space). Table 3 summarizes the results when generating these planograms using high and low customization levels. The percentage values are relative to the handmade version. GAP Generation is able to improve the manual planogram performance in all four metrics, both in the low and high customization versions. As expected, reducing the number of rules imposed on the planogram yields additional gains with an average increase of potential sales from 0.4% to 1.4%. We observed that the gains were more relevant in the metrics regarding the day-supply values (whose average and standard deviation were reduced by 38% and 61% respectively), which is consistent with our primary objective. Note that days-supply

values have a major impact on replenishment operations, holding costs and product availability. The planograms' filling rate also increases in both versions by 3% compared to the handmade planograms.

Table 3 To assess the impact of GAP we analyzed 25 Generation and Replication processes. Generation processes were run with low and high customization settings. Both processes bring benefits to the current practice, both in terms of potential sales increase and inventory management.

	GAP Generation		GAP Replication
	Low Customization	High Customization	
Potential sales increase *	1.4%	0.7%	–
Average days-supply reduction*	37.6%	35.3%	34.3%
Standard deviation days-supply reduction*	60.8%	51.4%	56.3%
Space Occupation**	96.5%	96.7%	97.3%
Execution Time (hh:mm:ss)	00:07:12	00:01:50	00:02:10

* with respect to the handmade planogram; ** 94% in handmade planogram

GAP Replication, by definition, has less margin for optimization, with changes being allowed only to the number of facings that products have in the new planograms, subject to the many allocation rules that were extracted from the role planogram. Nevertheless, a smarter product facing allocation can optimize the day-supply values. We used the same 25 examples to assess GAP Replication performance by replicating the handmade planograms to the same store. The results in Table 3 prove that we are still able to significantly improve days-supply balancing.

This project also lead to a significant cultural change in the company. The success of the project motivated the use of OR approaches (and more generically speaking, of analytical approaches) in other operations planning activities. In particular, it triggered many other projects with the same OR group from the University of Porto, both in space related problems, such as backroom optimization, and in other related areas, such as marketing, store operations and logistics.

Standardization: Knowledge Management for a Global Process

Space managers are divided into groups responsible for subsets of categories. The knowledge associated with each category is kept inside each group, supported by manuals that record the implementation details (using a template similar to Figure 2). Nevertheless, these manuals are frequently limited to the upper criteria levels, giving only a general idea of the reasoning behind the planograms. Consequently, most of the category knowledge resides with the space managers, which is partially lost when organizational changes occur. GAP also had a major impact on standardizing information and managing knowledge. First, the use of electronic generation manuals

(and the possibility of reusing them) centralized the categories' space planning knowledge and made it possible to systematize the tacit knowledge available into information that can be shared among peers. Second, it reduced the subjectivity of the process, which is now less dependent on the managers' experience.

Concluding Remarks

In the highly competitive retail environment of today, retailers can benefit from analytic tools for better decision making and many successful examples are reported in the literature. Shelf-space planning is one area that is still to be explored, mainly because of its complexity and high dependency on merchandising rules. We believe that this work is an important contribution in this direction, both from a theoretical and practical point of view. On the scientific front, we provide innovative mathematical models and efficient algorithms for the shelf-space allocation problem and bring more realism to the scientific approaches to this problem. From the application perspective, we give insights on how to tailor analytical approaches to the practice of shelf-space management, by introducing the replication problem and by allowing users to control the level of customization from solutions, while still applying optimization at every step of the process. For Sonae MC, this resulted in an automated process that still took advantage of the space managers' experience and knowledge. We also provide project management details that were critical during GAP implementation in Sonae MC, the major Portuguese retail company that partnered in this project.


Although this paper describes a real application of shelf-space planning, the approach does not intrinsically depend on any company specific policies, because it is based on rules that are defined at run-time. Therefore, it is sufficiently generic to be suitable for other retail companies working in the grocery industry or similar industries. Its modular nature also enables its adaptation and integration with other IT systems.

Appendix A: Target Facings Model

In this appendix, we provide a mathematical formulation of the Target Facings Model. Consider a specific category of a store with overall capacity C . The retailer wants to allocate N products, indexed by $i \in \mathcal{N}$, with length a_i . Each product is associated with a *space-to-sales* curve (Figure 5), which is linearized with piecewise lines, obtained using the facings associated with the days-supply intervals (Figure 11). There are T days-supply intervals, indexed by $n \in \mathcal{T}$. For each product i , the minimum and maximum facings of each interval n are ds_i^n and ds_i^{n+1} .

This *space-to-sales* curve represented in Figure 5 is widely used in the literature and is associated with a polynomial function depending on the space-elasticity parameter as first introduced by Corstjens and Doyle (1981). Other authors

Figure 11 The space elasticity curve is linearized using piecewise linearization based on the days-supply intervals represented by the vertical lines.



such as Gajjar and Adil (2010) have proposed piecewise linear approximations. However, we are the first to consider the problem with days-supply intervals.

The objective is to maximize the planogram's expected demand by determining the number of facings for each product i . The decisions to be made are: γ_n , which specifies whether the days-supply interval n is selected for all products; and W_i^n , which indicates the number of facings of product i if the days-supply interval is γ_n .

The formulation is as follows:

$$\text{Maximize } \sum_{i \in \mathcal{N}} \sum_{n \in \mathcal{T}} f_i^n(W_i^n) \quad (1)$$

$$\text{subject to: } \sum_{i \in \mathcal{N}} \sum_{n \in \mathcal{T}} W_i^n \cdot a_i \leq C \quad (\text{shelf-space capacity}) \quad (2)$$

$$l_i \leq \sum_{n \in \mathcal{T}} W_i^n \leq u_i, \forall i \in \mathcal{N} \quad (\text{minimum and maximum facings}) \quad (3)$$

$$ds_i^n \cdot \gamma_n \leq W_i^n \leq ds_i^{n+1} \cdot \gamma_{n+1}, \forall i \in \mathcal{N}, n \in \mathcal{T} \quad (\text{number of facings}) \quad (4)$$

$$\sum_{n \in \mathcal{T}} \gamma_n = 1 \quad (\text{single day-supply interval}) \quad (5)$$

$$W_i^n \in \mathbb{N}_0, \forall i \in \mathcal{N}, n \in \mathcal{T}; \gamma_n \in \{0, 1\}, \forall n \in \mathcal{T} \quad (\text{integrality}) \quad (6)$$

Appendix B: Generation Model

In this appendix, we provide a simplified mathematical formulation of the Generation Model. The notation used in Appendix A is extended by introducing K shelves, indexed by $k \in \mathcal{K}$ and with length w_k . The shelf attractiveness of Figure 6 is modeled by parameter β_k . Additionally, p_i is the profit per facing of product i and \bar{l}_i and \bar{u}_i are the lower and upper bounds for the number of facings of product i . These bounds are determined by the upstream Target Facings Model and correspond to the bounds of the selected days-supply interval ($\bar{l}_i = \sum_n \gamma_n \cdot ds_i^n$ and $\bar{u}_i = \sum_n \gamma_n \cdot ds_i^{n+1}$).

The objective is to maximize the planogram's expected profit by determining the number of facings for each product i and its placement on the shelves, while respecting all merchandising rules, translated into a tree diagram. Let W_{ik} be

the integer number of facings of product i on shelf k . The objective function and product allocation constraints can be defined as follows:

$$\text{Maximize } Z = \sum_{i \in \mathcal{N}} \sum_{k \in \mathcal{K}} p_i \cdot \beta_k \cdot W_{ik} \quad (7)$$

$$\text{subject to: } \bar{l}_i \leq \sum_{k \in \mathcal{K}} W_{ik} \leq \bar{u}_i, \quad \forall i \in \mathcal{N} \quad (\text{minimum and maximum number of facings}) \quad (8)$$

$$\sum_{i \in \mathcal{N}} a_i \cdot W_{ik} \leq w_k, \quad \forall k \in \mathcal{K} \quad (\text{shelf capacity}) \quad (9)$$

To ensure a correct positioning of the products within its product family and on the shelf, let $u \in \mathcal{M}$ denote the set of product families and \mathcal{N}_u the set of products of family u . To capture the tree diagram of Figure 7 consider that the set \mathcal{V}_u contains the immediately downstream families of parent family u , which for the final level corresponds to products. The family sequencing in the planogram is ensured by the multi-level network flow constraints presented below, where Y_{mk} (0/1) defines if family m (or product i) is located on shelf k and T_{mnk} (0/1) captures if family m is displayed immediately after family n on shelf k .

$$\sum_{m \in \mathcal{V}_u} T_{0mk} = Y_{uk}, \quad \forall u \in \mathcal{M}, k \in \mathcal{K} \quad (\text{first family belonging to parent family } u \text{ displayed in the shelf}) \quad (10)$$

$$\sum_{m \in \mathcal{V}_u} T_{m0k} = Y_{uk}, \quad \forall u \in \mathcal{M}, k \in \mathcal{K} \quad (\text{last family belonging to parent family } u \text{ displayed in the shelf}) \quad (11)$$

$$\sum_{n \in \mathcal{V}_u \cup \{0\}} T_{nmk} = Y_{mk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{only one family immediately before}) \quad (12)$$

$$\sum_{n \in \mathcal{V}_u \cup \{0\}} T_{mnk} = Y_{mk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{only one family immediately after}) \quad (13)$$

$$Y_{mk} \leq Y_{uk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{only on the shelf if the parent family is allocated}) \quad (14)$$

$$W_{ik} \geq Y_{ik}, \quad \forall i \in \mathcal{N}, k \in \mathcal{K} \quad (\text{only display facings on the shelf if the product is allocated}) \quad (15)$$

To complete the sequencing decisions and distribute the shelf length among the several product family layers as depicted in Figure 8 the following single commodity constraints are required:

$$\sum_{m \in \mathcal{V}_u} F_{0mk} = 0, \quad \forall u \in \mathcal{M}, k \in \mathcal{K} \quad (\text{null initial flow in each network}) \quad (16)$$

$$\sum_{m \in \mathcal{V}_u} F_{m0k} = \sum_{i \in \mathcal{N}_u} L_{ik}, \quad \forall u \in \mathcal{M}, k \in \mathcal{K} \quad (\text{the end flow is equal to the sum of the family product's length}) \quad (17)$$

$$\sum_{\substack{n \in \mathcal{V}_u \cup \{0\}: \\ m \neq n}} F_{nmk} + \sum_{i \in \mathcal{N}_m} L_{ik} = \sum_{\substack{n \in \mathcal{V}_u \cup \{0\}: \\ m \neq n}} F_{mnk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{flow conservation constraints}) \quad (18)$$

$$F_{mnk} \leq w_k \cdot T_{mnk}, \quad \forall u \in \mathcal{M}, m, n \in \mathcal{V}_u \cup \{0\} : m \neq n, k \in \mathcal{K} \quad (\text{flow can only traverse the selected arcs}) \quad (19)$$

$$\sum_{i \in \mathcal{N}} L_{ik} = w_k, \quad \forall k \in \mathcal{K} \quad (\text{all shelf length must be assigned}) \quad (20)$$

$$a_i \cdot W_{ik} \leq L_{ik}, \quad \forall i \in \mathcal{N}, k \in \mathcal{K} \quad (\text{the shelf length assigned must hold the facings to be displayed}) \quad (21)$$

where L_{ik} represents the shelf length assigned to product i on shelf k and F_{mnk} is the cumulative assigned shelf length (commodity flow) when changing from family m to family n on shelf k . This is used to describe the networks of Figure 8. **Note that constraints (20) and (21) replace constraints (9).**

The final set of constraints guarantee the rectangular shape of product families in the planogram. **We impose the rectangular shapes by determining the left and right coordinate of each family and ensuring that the family's first and**

last products of each shelf are kept within v units from these coordinates. To complete the rectangular shape, we force that the family has to be in consecutive shelves. To do so, let X_m^s (X_m^e) be the left (right) horizontal coordinate of the family m and FL_{mk} (LL_{mk}) be equal to one if k is the first (last) shelf of block m .

$$X_m^s \geq X_u^s + \sum_{\substack{n \in \mathcal{V}_u \cup \{0\}: \\ n \neq m}} F_{nmk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{left alignment}) \quad (22)$$

$$X_m^s \leq X_u^s + \sum_{\substack{n \in \mathcal{V}_u \cup \{0\}: \\ n \neq m}} F_{nmk} + C \cdot (1 - Y_{mk}), \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{left alignment}) \quad (23)$$

$$X_m^e - X_m^s \geq \sum_{i \in \mathcal{N}_m} L_{ik}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{right alignment}) \quad (24)$$

$$X_m^e - X_m^s - v \leq \sum_{i \in \mathcal{N}_m} L_{ik} + C \cdot (1 - Y_{mk}), \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} \quad (\text{right alignment}) \quad (25)$$

$$\sum_{k \in \mathcal{K}} FL_{mk} = 1, FL_{m0} = Y_{m0}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u \quad (\text{first shelf identification}) \quad (26)$$

$$\sum_{k \in \mathcal{K}} LL_{mk} = 1, LL_{mK} = Y_{mK}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u \quad (\text{last shelf identification}) \quad (27)$$

$$FL_{m,k+1} + Y_{mk} = Y_{m,k+1} + LL_{mk}, \quad \forall u \in \mathcal{M}, m \in \mathcal{V}_u, k \in \mathcal{K} : k \neq K \quad (\text{shelf consecutiveness}) \quad (28)$$

Additional details about the generation model are presented in Bianchi-Aguiar et al. (2015b).

Appendix C: Replication Model

In this appendix, we provide a simplified mathematical formulation of the Replication Model. We will follow the notation used in Appendices A and B.

The objective is to replicate the product placement of a given planogram to a new store without the need to specify merchandising rules, while maximizing the new planogram's expected profit. For that, products should be assigned by the same sequence to the same shelves and product families should keep their rectangular shapes, which can be ensured by aligning the family's first and last product of each shelf (with a tolerance of v units). For that, consider the sets \mathcal{N}_k containing the products from each shelf k and M^L left-alignments and M^R right-alignments, indexed by $m \in \mathcal{M}^R \cup \mathcal{M}^L$. As products are not able to change between shelves, the shelf attractiveness was incorporated in the profit: \bar{p}_i .

Let the products be numbered according to their positioning. W_i is the integer number of facings of the i^{th} product and L_i is the shelf length assigned to the i^{th} product i . Additionally, X_m^L and X_m^R capture the horizontal location of left and right alignments, respectively. The replication model reads:

$$\text{Maximize } Z = \sum_{i \in \mathcal{N}} \bar{p}_i \cdot W_i \quad (29)$$

$$\text{subject to: } \bar{l}_i \leq W_i \leq \bar{u}_i, \quad \forall i \in \mathcal{N} \quad (\text{minimum and maximum number of facings}) \quad (30)$$

$$L_i - a_i W_i \geq 0, \quad \forall i \in \mathcal{N} \quad (\text{the shelf length assigned must hold the facings to be displayed}) \quad (31)$$

$$\sum_{i \in \mathcal{N}_k} L_i = w_k, \quad \forall k \in \mathcal{K} \quad (\text{all shelf length must be assigned}) \quad (32)$$

$$X_i = \sum_{j \in \mathcal{N}_k: j < i} L_j, \quad \forall k \in \mathcal{K}, i \in \mathcal{N}_k \quad (\text{product's positioning and relationship}) \quad (33)$$

$$X_m^L - v \leq X_i^s, \quad \forall m \in \mathcal{M}^L, i \in \mathcal{N}_m^L \quad (\text{product families' left alignment}) \quad (34)$$

$$X_m^L \geq X_i^s, \quad \forall m \in \mathcal{M}^L, i \in \mathcal{N}_m^L \quad (\text{product families' left alignment}) \quad (35)$$

$$X_m^R + v \geq X_i^s + L_i, \quad \forall m \in \mathcal{M}^R, i \in \mathcal{N}_k^R \quad (\text{product families' right alignment}) \quad (36)$$

$$X_m^R \leq X_i^s + L_i, \quad \forall m \in \mathcal{M}^R, i \in \mathcal{N}_k^R \quad (\text{product families' right alignment}) \quad (37)$$

This formulation requires each product to stay in a single product. If this is not the case, then each products' position should be assigned with a diferent index and small changes are needed in the formulation. Additional details about the replication model are presented in Bianchi-Aguiar et al. (2015a)

Acknowledgments

The authors are grateful to other members of the team. We give a special thanks to the three space managers, Constantino Gomes, Pedro Soares and Susana Borges; and also Frederico Santos, Joel Pacheco, Miguel Camanho and Helder Matos from ISI. We would also like to acknowledge the relevant comments from the 2015 Daniel Wagner Prize judges, with a special thank to Ann Bixby, and also to J.H.Tice, for the English revision.

The authors are also grateful to the FCT – Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) – for awarding the grants SFRH / BD / 74387 / 2010 and SFRH / BPD / 98981 / 2013. This work is also financed by the ERDF – European Regional Development Fund – through the ON.2 Programme, and by National Funds through the FCT within Smart Manufacturing and Logistics [Project NORTE - 07 - 0124 - FEDER - 000057].

References

- Bai, Ruibin. 2005. An Investigation of Novel Approaches For Optimising Retail Shelf Space Allocation. PhD Thesis. *The University of Nottingham*.
- Baker, RC, Timothy L Urban. 1988. A deterministic inventory system with an inventory-level-dependent demand rate. *Journal of the Operational Research Society* **39**(9) 823–831.
- Bianchi-Aguiar, Teresa, Maria Antónia Carravilla, José F. Oliveira. 2015a. Replicating shelf space allocation solutions across retail stores. *Under revision on a peer reviewed journal*.
- Bianchi-Aguiar, Teresa, Elsa Silva, Luis Guimar aes, Maria Antónia Carravilla, José F. Oliveira. 2015b. Allocating products on shelves under merchandising rules: multi-level product families with display directions. *Under revision on a peer reviewed journal*.
- Chandon, Pierre, J Wesley Hutchinson, Eric T Bradlow, Scott H Young. 2009. Does In-Store Marketing Work ? Effects of the Number and Position of Shelf Facings on Brand Attention. *Journal of Marketing* **73**(6) 1 – 17.
- Corstjens, Marcel, Peter Doyle. 1981. A Model for Optimizing Retail Space Allocations. *Management Science* **27**(7) 822–833.
- Curhan, Ronald C. 1972. The Relationship Between Shelf Space and Unit Sales in Supermarkets. *Journal of Maketing Research* **9**(4) 406–412.
- Drèze, Xavier, Stephen J. Hoch, Mary E. Purk. 1994. Shelf management and space elasticity. *Journal of Retailing* **70**(4) 301 – 326.

- Gajjar, Hasmukh, Gajendra Adil. 2010. A piecewise linearization for retail shelf space allocation problem and a local search heuristic. *Annals of Operations Research* **179**(1) 149–167.
- Geismar, H. Neil, Milind Dawande, B.P.S. Murthi, Chelliah Sriskandarajah. 2014. Maximizing revenue through two-dimensional shelf-space allocation. *Production and Operations Management* Available online.
- Hübner, Alexander H., Heinrich Kuhn. 2012. Retail category management: State-of-the-art review of quantitative research and software applications in assortment and shelf space management. *Omega* **40**(2) 199 – 209.
- Keltz, Heather, Kevin Sternecker. 2009. The trend toward consumer–centric merchandising requires assortment management and space planning investments. Tech. rep., AMR Research.
- Kurtulus, Mümin, L. Beril Toktay. 2009. Category captainship practices in the retail industry. *Retail Supply Chain Management: Quantitative Models and Empirical Studies*. (Springer, New York), 79–98.
- Pieters, Rik, Michel Wedel, Rajeev Batra. 2010. The Stopping Power of Advertising: Measures and Effects of Visual Complexity. *Journal of Marketing* **74**(5) 48–60.
- Russell, Robert A., Timothy L. Urban. 2010. The location and allocation of products and product families on retail shelves. *Annals of Operations Research* **179**(1) 131–147.
- Urban, Timothy L. 1998. An inventory-theoretic approach to product assortment and shelf-space allocation. *Journal of Retailing* **74**(1) 15 – 35.
- Yang, Ming-Hsien, Wen-Cher Chen. 1999. A study on shelf space allocation and management. *International Journal of Production Economics* **61**(510) 309–317.